RadNet Protocol Guidelines
Thursday, February 19, 2004

RadNet Protocol Guidelines

1. The RadNet protocol is the format of the data contained within a UDP data-gram. The structure of UDP messages is outlined in RFC's 768,862-865,867, and 1119. Further information is contained in the reference section of this document.

2. The RadNet protocol is split into three functional areas: the header, body, and footer. These three areas are not evident when examining a RadNet message, and are not interdependent. The intent of this format is to provide the most basic information (which instrument and its status) first, in a simple and consistent manner to facilitate programming. The header format is the same for all instruments and contains basic information about the instrument such as its location, type, status, and address. If the instrument is operating properly, or simple programming is desired there may be no need to collect (process and store) additional information. Detailed information specific to an instrument's settings and readings is contained in the message body. Alarm settings and readings may be sufficient knowledge, or the detailed information from individual channels may be desired. The footer contains the number of channels, and can repeat its format for each of those channels. Information not otherwise supported can be placed in a single footer.

3. Data may be added at the end of the footer. The RadNet protocol is a known length, thus information beyond the specified length is not supported by the protocol, but can nonetheless be used by other programs. This practice is not advisable as the protocol may have features added in the future. Using a pass-through message type that RadNet supports would be more prudent. This message type is described in greater detail in the Header section of this document.

4. At this time RadNet does not address remote calibration, changing monitor settings, or changing configurations remotely. These functions may be addressed in the future. These options can be supported with the pass-through message. RadNet only specifies the format of the first 3 bytes of the pass-through message. The instrument vendor must provide the remaining specifications, as it is anticipated that these functions will be unique to each vendor.

5. Backward compatibility must be provided at all times with client/server software. New versions of RadNet will not hinder previous versions of supporting software.

6. Every RadNet message has a header. Transactional instruments (such as PCM's,) will have a body and possibly footer for transactions, but may only have a header for (normal) status messages. Non-transactional instruments (such as area monitors) include a body or footer. Transactions will be pushed whenever a transaction has occurred. Status changes will result in a push as well.

7. All RadNet byte and word values are unsigned. Word values follow the Intel 8088 CPU format of least significant byte (LSB), most significant byte (MSB). To obtain the value for a word field use the following formula: (MSB * 255) + LSB = 16 bit integer value. All char values are 8 bit standard ASCII coded values. All Float values are standard IEEE 32-bit (4-byte) floating-point numbers.

8. If a vendor needs a new status code, monitor type code or any other addition to the RadNet protocol, contact the protocol working group Chairperson, (http://www.rad-net.net) and request the addition be made to the RadNet protocol code list. Assigning a non-listed code may render a device non-compliant with the standard. Codes may be assigned and not be listed between protocol versions.

9. If an instrument does not support the RadNet value, the value must be padded to the proper length. For example, if an instrument (e.g. PCM) does not support RWP entry (a character value), the RWP bytes must be shipped as spaces (ASCII 32). Number values must be shipped as a zero value. For example, if operational status (Header byte 9) is not supported, it must be pushed as a zero value "00000000b".

10. Codes are listed in Appendix A. Codes do not change between instrument types or location in the message (e.g. Operational Status, which is used for all instruments, and in the header and footer).

11. If supported by the instrument manufacturer, RadNet can support a polling methodology. This polling is accomplished by setting the instrument "abnormal" and "normal" push rate values to 0 (zero) . This setting will tell the instrument to stop sending out data packets. The client/monitoring computer then sends a RadNet Request Data (Code=5) message. The

instrument must be capable of receiving and understanding a RadNet. When the instrument receives this message, it will then ship the packet using a standard data packet. When using the polling methodology, care must be taken to reduce network errors. Since most instruments are simple 8 or 16 bit processors incapable of spawning multiple threads, a "Destination Port Unreachable" error may occur when using a multiple polling computer. This error increases with the number of polling computers and polling frequencies. Polling will have a dramatic effect upon the bandwidth usage for the network and should be monitored by your network team.

12. If supported by the instrument manufacturer, two-way communication can be implemented using the following methods: RadNet has two UDP (preferred IP protocol) and TCP ports assigned for its use. These ports are 16367 and 16368. Port 16367 can be used to push (send) data onto the network for monitoring. Port 16368 subsequently be used to receive request/commands from monitoring/client computers. This approach allows the instrument to continue pushing data while receiving commands/requests. If more than one computer on the network is sending requests, care must be taken to prevent the "Destination Port Unreachable" network error. Commands can be received and data pushed using the same port (such as 16367), however, these sending and receiving functions can not occur simultaneously using the same port. If the instrument is sending and receiving on the same port, then the sending computer must handle the "Destination Port Unreachable" network error. This error will increases with the number of computers sending requests and is also affected by the push rate of the instrument.

13. RadNet supports sending broadcast messages to an instrument. Sending Broadcast request/command messages to instruments using the following methodology: A broadcast message is described as all computers on a subnet receiving the same packet of data. A more detailed description can be found in the UDP protocol RFC. If the sending computer wishes to tell all instruments on a subnet to perform a diagnostic, then it sends out .255 (class c = x.x.x.255 or class b = x.x.255.255) to that subnet. The broadcast value may be different if the network team has subneted the subnet on the your network. Users should always check with the network team to obtain the broadcast value. When the instrument receives the broadcast message, it then performs the requested task. The monitoring computer is able track who performed the task by comparing the learned list of instruments to who responded to the request. This interaction is one simple way of making sure all instruments are still online.

14. Monitoring computers learn what instruments are on the network and detect their data frequency using the following methodology: When an instrument pushes its first packet of data, the monitoring/client computer learns about the instrument (captures its IP address and timestamps the packet). When the monitoring computer receives the next packet, the monitoring computer can calculate the instrument's push frequency. With these two data points, the monitoring computer can then detect when an instrument goes off-line or the network has gone down. The monitoring computer has one of two options: first, to do nothing, second, to take actions. These actions may include Pinging the instrument, or sending a RadNet message (request data, code =5). If no response is received from the instrument, then the monitoring computer can page and email that the instrument is no longer on the network.

15. Filtering RadNet data using different ports: RadNet can filter different instrument types by assigning a different port to each instrument type. This approach is very effective in a large instrument deployment. For example: Alpha CAM's can use port 16367, Beta CAM's 16368, Area Monitors 16369, and PCM can be assigned to port 16370. The client/monitoring computer software will then toggle between the ports and only instruments on that particular port would be seen by the computer.

16. Additionally, it is possible to filter instruments in areas by assigning a different port to each area. All instruments in one room can be assigned to port 16367 and another room can be assigned to port 16368. By toggling the ports on the monitoring computer, only the instruments for that room would thus be seen. When using this approach, it is important to have sniffed the network to ensure no other devices are using the port.

17. RadNet supports the three-tier architecture using the following methodology: RadNet supports the use of broadcasting messages onto the network. Using this functionality allows the end user to implement a three tier architecture. When a broadcast message is sent (pushed) onto to the network, all computers on the subnet will receive identical packets of information. Each middle tier computer can then handle the packet in different ways. One computer might be tasked with archiving the packet to a database, while another handled paging and email, and still another supplied a web interface.

18. It is important to remember that a single computer can handle all of the tasks outline above. If, however, at some point the work load increases, then other computer could be added to reduce the work load of the monitoring computer. One major concern to using a single monitoring computer is the idea of "the single point of failure". If different computers are deployed

to manage a specific task and a failure should occur within one task, only its associated computer would also fail, leaving the others to continue operating.

19. RadNet supports Redundancy using the following methodology: Redundancy can be achieved by placing one to x number of monitoring computers onto the network. A "master" computer can receive and process RadNet. The master computer must send a status message to a "slave" computer (backup computer). As long as the master computer continues to send messages to the slave, that computer continues in the capacity of slave. If the master computer fails to send a message to the slave computer, the slave computer automatically becomes the master, functioning in the capacity of master rather than backup. When the original master computer resumes sending messages to the original slave computer, their original roles and functions will resume.

20. By placing the two computers on two or more subnets or behind two or more routers, network redundancy can be achieved. When this approach is used, the instruments must push data (broadcast) onto each subnet or across each router.

21. Network tuning is supported by RadNet using the following methodology: Because each instrument has its own push rate, the data can be push based upon the instrument/facility requirements. Instruments can thus be configured to push whenever they have valid data. If an Alpha CAM takes 10 minutes to produce new data, then the push rate should be set to 10 minutes, rather than 1 minute. This setting reduces the number of packets of data being pushed onto the network. Another way to reduce network traffic is to analyze the hazards within the area. If the work area has a low risk (such as a storage building), then data may only need pushing once daily. In other areas, a higher push rate may be implemented (such as every 1 minute) based upon the risk and past history. One of RadNet's fundamental rules is that an instrument shall push data upon any status change. Even though an instrument only normally pushes (sends) data once a day, if the instrument goes into an alarm condition, the instrument will then push a data packet at the abnormal push rate. This default provides real-time data acquisition for the alarm/trouble condition. The monitoring computer can also ping instruments at some frequency to make sure the instruments are still on the network.

    It is important to work with the network support team to monitor the effects of placing instruments onto the network. Before placing any instruments on the network, the network team should monitor the bandwidth usage over a given time period. This time frame should be long enough to capture the peak usage on the network. After completing this task instruments are placed on the network at the push rates that are needed by the room/area/facility. The network team should again monitor the network for bandwidth usage. The goal is to have enough bandwidth to support the instrument going into abnormal push rates. Only the instrument in the affected area will be sending data at the abnormal push rate.

22. Minimizing data loss on the network can be achieved by testing the network. Testing is easily done by taking the push rate of the instrument and comparing it to the number of packets received by the monitoring computer. Packet reliability should be around 99% for a typical network. If this is not the case, the network team should troubleshoot the network. One solution is to increase the push rate to ensure packet throughput. This increase must be done with caution! Before increasing the rate, install a network analyzer on the network to see if there is a device causing the packet loss. If the network bandwidth is consistently grater than or equal to 60%, then increasing the push rates may make data loss worse. On a typical Ethernet network, increasing the push rate to decrease packet loss could increase network collision thus having the opposite of the desired effect. Again, the only way to see if this is happening is to connect a network analyzer and monitor the health of the network.

23. Instrument manufacturers should provide a means to troubleshoot RadNet communication from the field. Supplying a simple PING routine within the hardware/software can significantly increase the ability to diagnose networking problems. The field person is then able to test the network connection from the instrument back to the monitoring computer. If the device is a serial-to-Ethernet converter, then a method to monitor the serial and network data stream should also be provided.